

Secrets of Text Anti-Aliasing: Revealed!

John Platt, Sr. Researcher, Statistical Media Processing
Turner Whitted, Sr. Researcher, Hardware/Graphics

April 24, 2003


Work smarter. Get training.

<http://mste>

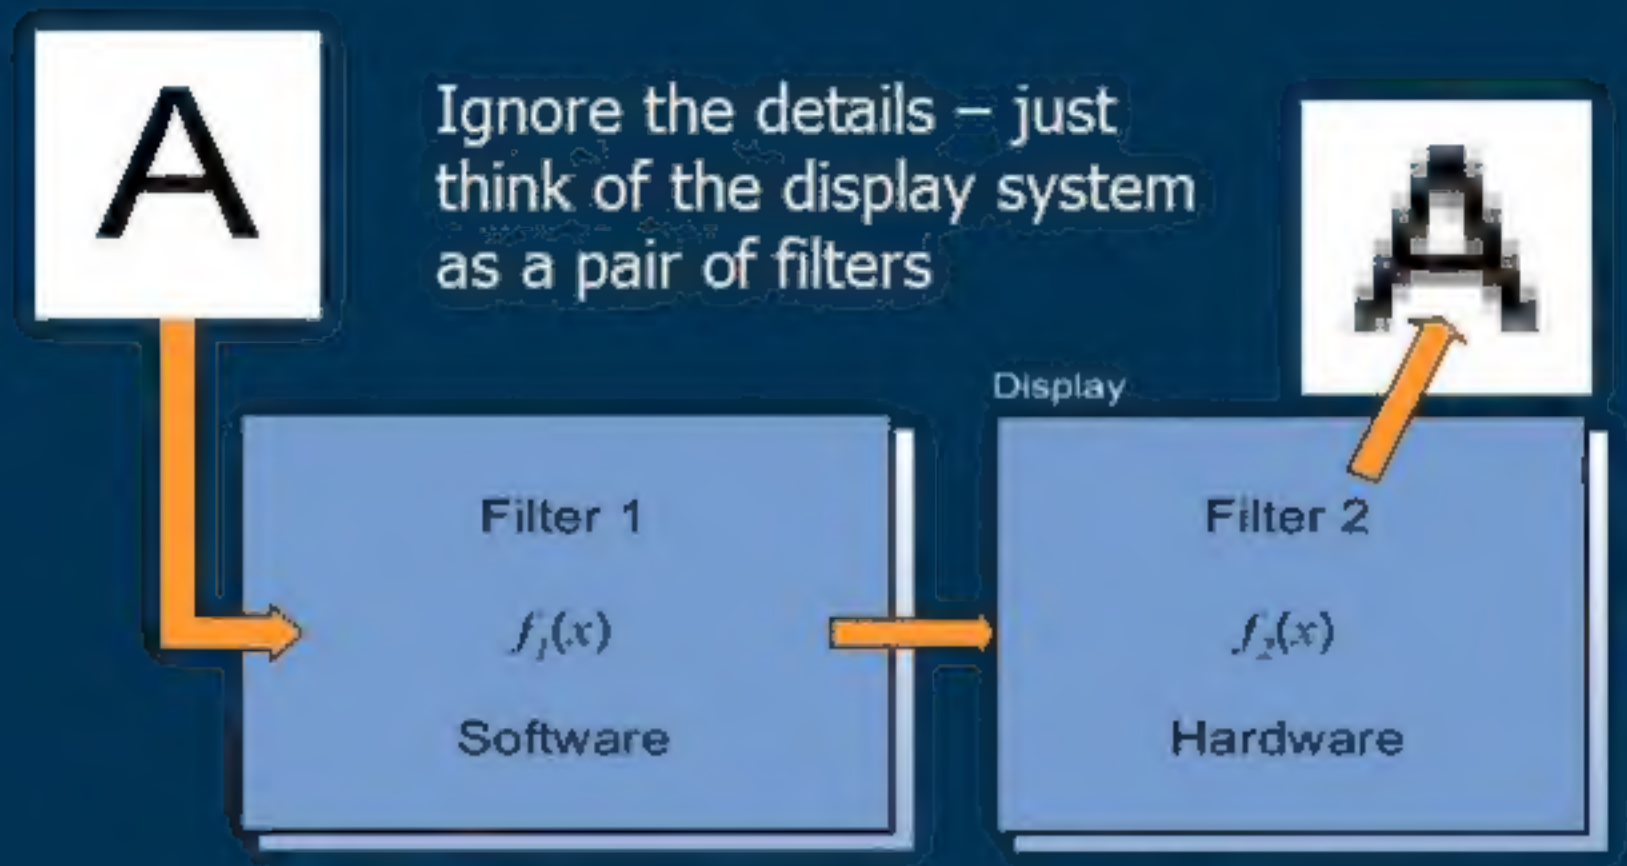
Structure of Talk

- Two sub-talks
 1. Signal Processing Background (Turner)
 - Linear systems theory
 - Frequency domain explanation of aliasing
 - Debunk misconceptions about anti-aliasing
 2. Text Anti-Aliasing (John)
 - Text-Specific Anti-Aliasing (Visio mode)
 - Display-Specific Anti-Aliasing (ClearType)

Display System Overview




Display System Overview



What's a Filter?

Response

Example: low pass filter (LPF), $f(x)$



What's a Filter?

Convolution

Input signal



Filter impulse response



Filtered output signal




What's a Filter?

Properties of linear systems


- Convolution
 - The output of a linear system is the convolution of the input with the system response
 - Response of cascaded linear elements is the convolution of the individual responses
- Dual representation
 - Spatial domain and frequency domain representations of a signal are related by the Fourier transform
 - Convolution in the spatial domain is equivalent to multiplication in the frequency domain

Signals: Superposition

$\text{Signal 1} - 0.5 \text{Signal 2} + 0.3 \text{Signal 3} = \text{Resulting Signal}$



Signals: Dual Representation




Signals: Sampling and Reconstruction

Densely sampled low frequency signal



Frequency domain representation of sampled signal



How to reconstruct the original continuous signal?

Signals: Aliasing

Higher frequency signal



Frequency domain representation of sampled signal



Signals: Aliasing

Still higher frequency signal



Frequency domain representation of sampled signal



Replicated spectra overlap

Signals: Unlimited Frequency

Cross section of letter "O" in center of image



Spectrum has infinite extent

don't even think about

sampling without

aliases




Filtering: Anti-aliasing

Spatial domain

Frequency domain

convolve

multiply



Filtering: Implementation

• How does one filter a digital signal without first sampling?

- Oversample

- Sample at high rate
- Perform discrete convolution
- Decimate at display resolution


- Analytic

- In special cases, a closed form solution to the convolution integral is available

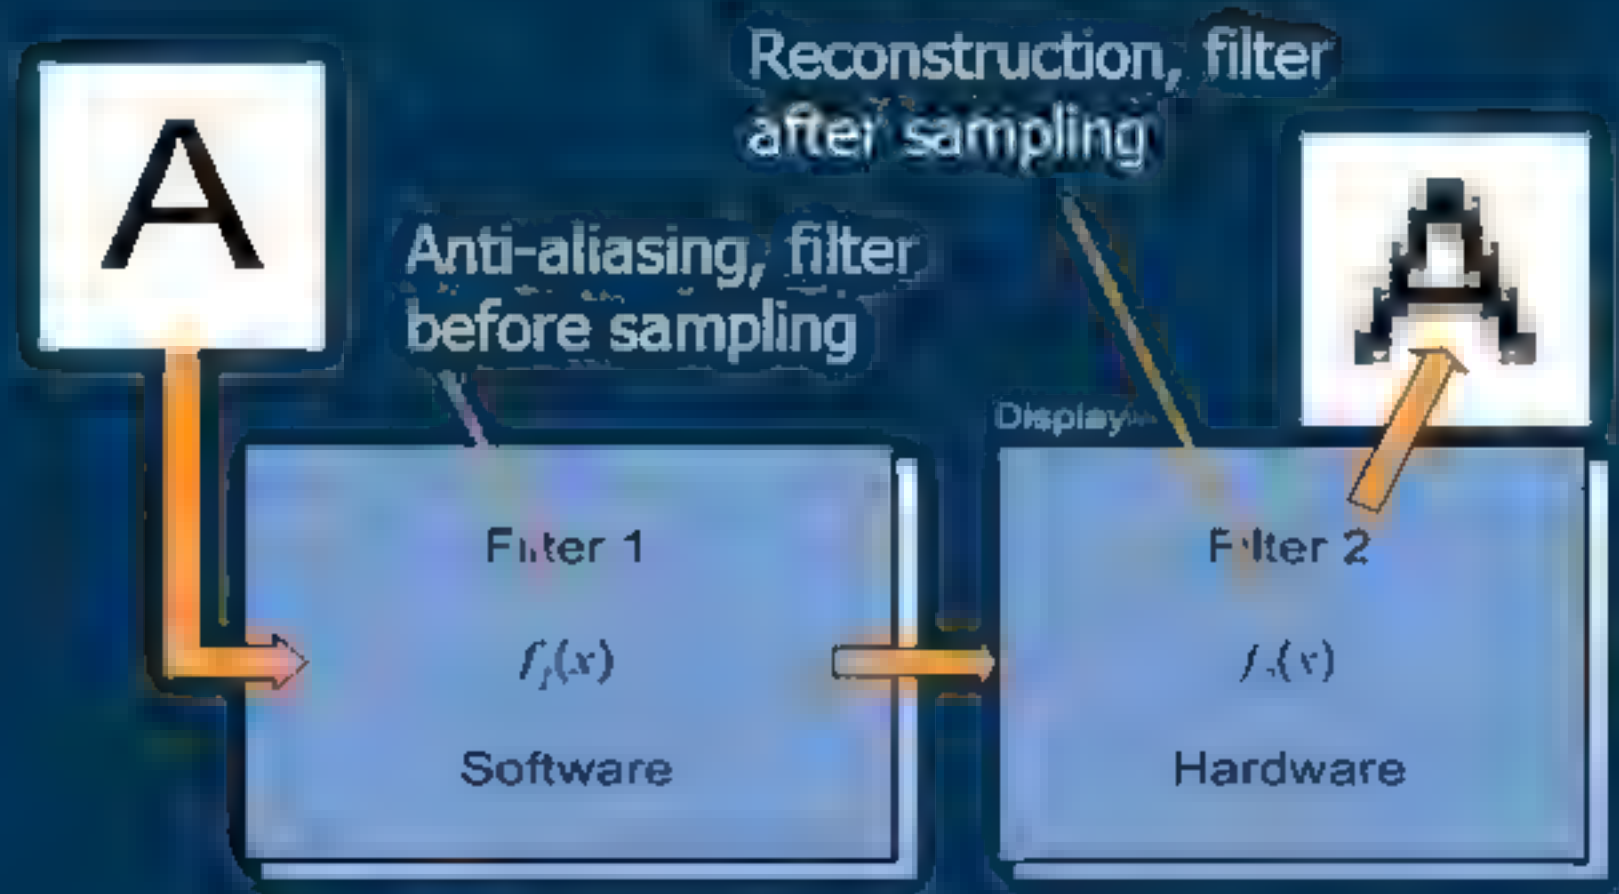
Filtering: Reconstruction in Display Hardware



Frequency response of display chain



Display System Review



Text Anti-Aliasing

• Text-Specific Anti-Aliasing

- text display is different from generic image display

• Display-Specific Anti-Aliasing

- use spatial structure of LCDs to gain more resolution


Text Specific Anti-Aliasing

- Text is different from generic images
 - comes as glyphs: pre-defined objects with meaning
 - defined as foreground/background blend
 - often occupies entire dynamic range of color
- Generic text anti-aliasing
 - specialized to adapt to differences, above
- Microsoft-specific tricks
 - fix up some problems with generic text anti-aliasing

Font Display

- Convert ideal font glyph into pixels

a



Steps of Font Display (in Software)

Hinting of Font Outlines (optional)

Sample Outlines at High Resolution

Filter Samples down to Pixel Resolution

Composite using Foreground and Background Colors

Correct for Display Non-linearity

Text is Defined as an Alpha Mask

- Glyphs have no innate color, just opaqueness (alpha, α)

$\alpha = 1$ inside glyph



m

$\alpha = 0$ outside glyph

- Compute pixel colors through *compositing*
 - $\text{color} = \alpha * (\text{foreground color}) + (1 - \alpha) * (\text{background color})$
- Fully correct steps
 - Composite, then anti-alias three color images
- To speed things up by a factor of 3:
 - filter the alpha image, *then* composite
 - perfectly correct if foreground & background are constant

Negative Lobes Do Not Work for Text

Convolve a 1-pixel wide stem



with a good anti-aliasing filter

get undisplayable
negative pixel values


Negative Lobes Cause Ringing




We are stuck with filters with positive weights only

First Impulse: Use a Box Filter

- Compute fraction of pixel covered by glyph
 - = one-pixel wide box filter sampled at pixel centers
- Resist this temptation





box filter



frequency response

Better: Single Lobe of a Sinc Filter

- First lobe of sinc filter, width = $3/2$ pixel



less aliasing

Two Dimensions: Two Filters


- Filter theory extends to two dimensions
 - 2D filters have 2D frequency responses
 - Sine waves with frequency & direction
 - 2D filters are computationally expensive (in software)
- Approximation: separable filters
 - First filter in x, then y (or visa versa)
 - Only need to think about 1D filters & responses

Positive-Only Filters Cause Lots of Gray

- Weighted average of white & black = gray!
- Positive-only filters tend to wash characters out
 - foreground colors get less intense

Use Non-Linearity to Combat Gray

- Gently push filtered α values towards foreground



- Gentle non-linearity: hard to analyze with filter theory
 - does not seem to affect aliasing too strongly
 - definitely helps with wash-out

Text Anti-Aliasing Results

1-abcdefghijklmnopqrstuvwxyz
2-abcdefghijklmnopqrstuvwxyz
3-abcdefghijklmnopqrstuvwxyz
4-abcdefghijklmnopqrstuvwxyz
5-abcdefghijklmnopqrstuvwxyz
6-abcdefghijklmnopqrstuvwxyz
7-abcdefghijklmnopqrstuvwxyz
8-abcdefghijklmnopqrstuvwxyz
9-abcdefghijklmnopqrstuvwxyz
10-abcdefghijklmnopqrstuvwxyz
11-abcdefghijklmnopqrstuvwxyz
12-abcdefghijklmnopqrstuvwxyz
13-abcdefghijklmnopqrstuvwxyz
14-abcdefghijklmnopqrstuvwxyz
15-abcdefghijklmnopqrstuvwxyz
16-abcdefghijklmnopqrstuvwxyz
17-abcdefghijklmnopqrstuvwxyz
18-abcdefghijklmnopqrstuvwxyz
19-abcdefghijklmnopqrstuvwxyz
20-abcdefghijklmnopqrstuvwxyz
21-abcdefghijklmnopqrstuvwxyz
22-abcdefghijklmnopqrstuvwxyz

box filtering
(shipped in Win2K)

1-abcdefghijklmnopqrstuvwxyz
2-abcdefghijklmnopqrstuvwxyz
3-abcdefghijklmnopqrstuvwxyz
4-abcdefghijklmnopqrstuvwxyz
5-abcdefghijklmnopqrstuvwxyz
6-abcdefghijklmnopqrstuvwxyz
7-abcdefghijklmnopqrstuvwxyz
8-abcdefghijklmnopqrstuvwxyz
9-abcdefghijklmnopqrstuvwxyz
10-abcdefghijklmnopqrstuvwxyz
11-abcdefghijklmnopqrstuvwxyz
12-abcdefghijklmnopqrstuvwxyz
13-abcdefghijklmnopqrstuvwxyz
14-abcdefghijklmnopqrstuvwxyz
15-abcdefghijklmnopqrstuvwxyz
16-abcdefghijklmnopqrstuvwxyz
17-abcdefghijklmnopqrstuvwxyz
18-abcdefghijklmnopqrstuvwxyz
19-abcdefghijklmnopqrstuvwxyz
20-abcdefghijklmnopqrstuvwxyz
21-abcdefghijklmnopqrstuvwxyz
22-abcdefghijklmnopqrstuvwxyz

improved filtering
(shipped in WinXP)



Text-Specific Anti-Aliasing Summary

- Goals are same as generic anti-aliasing
 - get rid of aliasing & blur, but don't spend much CPU
- Cannot use negative lobe filters
 - causes ringing and undisplayable negative pixels
- There are better filters than box filters
- Can use gentle non-linearity to combat font wash-out

ClearType

Original ClearType work with:

Claude Betrisey

Jim Blinn

Bodin Dresevic

Bill Hill

Greg Hitchcock

Bert Keely


Don P. Mitchell

Turner Whitted

Lots of other people since then!

Extra Resolution on LCDs?


- Three independent light sources per pixel
 - Three times as much resolution?
 - Only if you are color blind



- ClearType: trade off resolution for color error


Recall: Generic Anti-Aliasing

Let's reconsider anti-aliasing on all three color channels



Problem with Anti-aliasing on LCDs


- Standard anti-aliasing assumes pixels have no internal structure



Standard Anti-Aliasing: Position Error

- Position of character depends on color!


Red object shifted
 $\frac{1}{3}$ pixel to the left



Blue object shifted
 $\frac{1}{3}$ pixel to the right

True object location


- Black/white object appear at all three locations
 - These objects appear blurry



Displaced Sampling

Generate a sample for each sub-pixel

- at the location of the sub-pixel
- does not generate position error



New Problem: Color Aliasing

1 cycle/pixel black/white sine wave



yields uniform (DC) color!


Displaced sampling by itself does not work: need pre-filtering

just like normal anti-aliasing, except with colors

RGB Decimation


- Pre-filter to get rid of color aliasing
 - suppress frequencies at 2x pixel Nyquist (1 cycle/pixel)
 - keep frequencies above pixel Nyquist (extra resolution)
- Followed by displaced sampling
 - sample filter output at sub-pixel location
 - not at center of pixel

RGB Decimation



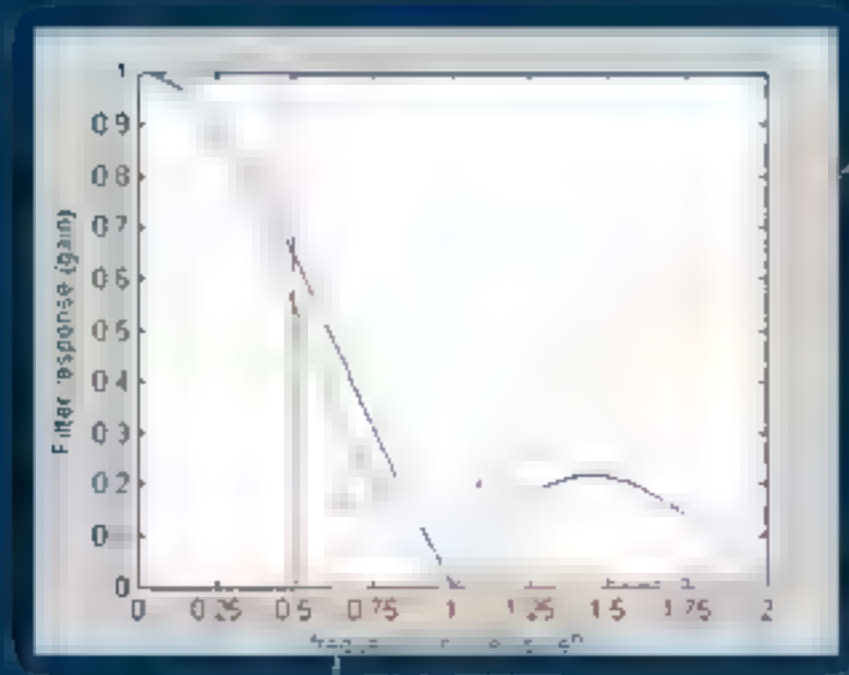
ClearType Filtering is Applied to Alpha

- Same trick as in standard text anti-aliasing



Which Filters to Use?

- Either 1-pixel wide  or 1.63-pixel wide single 



Good ClearType filters block out near 1 cycle/pixel


and leave extra resolution above pixel Nyquist

ClearType in Two Dimensions

- Use separable filters again (on α channel)
 - Parallel to stripes (typically vertical)
 - just use standard anti-aliasing
 - use 1.5 pixel one lobe sinc filter parallel to stripes
 - Perpendicular to stripes (typically horizontally)
 - use box filter or one lobe sinc filter perpendicular
 - use displaced sampling to get extra resolution
- Non-linear contrast enhancement works with ClearType

Results from ClearType Software

- horizontal filtering only:



no filtering

standard anti-aliasing

ClearType

Results Made More Obvious

- LCD sub-pixel intensities (shown in black/white)





no filtering

standard anti-aliasing

ClearType

Gamma Correction

- Gamma Correction compensates for display non-linearity



apply inverse function first to cancel out non-linearity

Gamma Correction is Vital for ClearType

- Display non-linearity \approx makes power at double frequencies

acceptable sine wave at 0.5 cycles/pixel

display non-linearity


sine wave at 1 cycle/pixel

color pattern of display


constant color!

- Gamma correction prevents color fringes

Gamma Correction Results



Gamma corrected



Not Gamma corrected

Summary of ClearType Filtering

- Use known positions of LCD sub-pixels to improve resolution
- Use RGB Decimation
 - Prefiltering to get rid of color aliasing
 - Displaced sampling to get more resolution
- Gamma Correction is Important

Summary of Talk

- Filtering can be understood by filter theory
 - Signals and systems are easy with frequencies
 - Filter before you sample to get rid of aliasing (jaggies)
- Text anti-aliasing is different from generic anti-aliasing
 - cannot use negative lobes
 - can use better filters than box filter
 - use non-linearity to compensate for wash-out
- ClearType is display-specific anti-aliasing
 - filtering + displaced sampling to get rid of color aliasing